A Parametric Factor Model of the Term Structure of Mortality

Niels Haldrup*
Siem Jan Koopman†
Carsten P. T. Rosenskjold*

January 31, 2017

Abstract

The prototypical Lee-Carter mortality model is characterized by a single common time factor that loads differently across age groups. In this paper we propose a factor model for the term structure of mortality where multiple factors are designed to influence the age groups differently via parametric loading functions. We identify four different factors: a factor common for all age groups, factors for infant and adult mortality, and a factor for the accident hump that primarily affects mortality of relatively young adults and late teenagers. Since the factors are identified via restrictions on the loading functions the factors are not designed to be orthogonal but can be dependent and can possibly cointegrate when the factors have unit roots. We suggest two estimation procedures similar to the estimation of the dynamic Nelson-Siegel term structure model. First, a two-step non-linear least squares procedure based on cross-section regressions together with a separate model to estimate the dynamics of the factors. Second, we suggest a fully specified model estimated by maximum likelihood via the Kalman filter recursions after the model is put on state space form. We demonstrate the methodology for US and French mortality data. We find that the model provides a good fit of the relevant factors and in a forecast comparison with a range of benchmark models we find that especially for longer horizons, variants of the paramtric factor model have excellent forecast performance.

*Center for Research in Econometric Analysis of Time Series (CREATES), Department of Economics and Business Economics, Aarhus University, Fuglesangs All 4, DK-8210 Aarhus V, Denmark. Corresponding e-mail: cthillemann@econ.au.dk

†Faculty of Economics and Business Administration, VU University Amsterdam, Tinbergen Institute, the Netherlands.

1 Introduction

The Lee and Carter (1992) (LC) model has become a benchmark model when estimating and forecasting improvements in age-specific death rates and the calculation of life expectancy. The model is basically a one-factor model that allows for a single common time trend with age-specific loadings. The model has been extended in many different ways. For instance Booth et al. (2002) and Renshaw and Haberman (2003) extend the model with a second common time trend with age-specific loadings. Hyndman et al. (2007) developed a functional data approach in which the data are smoothed across age prior to modelling using penalised regression splines and principal component analysis. De Jong and Tickle (2006) use the state space framework to build in smoothness in the LC model using b-splines. Czado et al. (2005) used Bayesian estimation to impose smoothness.

Typically, the estimation of factors is implemented via either singular value decomposition (SVD) or principal component analysis (PCA). For models with multiple factors, these are identified by making the factors orthogonal. Subsequently, the factors are modelled as individual time series models which can be used for forecast projections.

The LC model and its extensions are basically statistical models that summarize the variability of the measured age-specific death rates over time in a parsimonious way. No structure is imposed in the model specification. However, in the demographics literature it is well established that age groups are exposed rather differently to death risk and it seems reasonable to believe that separate time factors may affect these different age groups rather than just a single common factor as suggested by the basic Lee and Carter (1992) model.

Mortality laws for death rates observed at a given time have been suggested by amongst others Gompertz (1825), Makeham (1860) and Heligman and Pollard (1980); Tabeau et al. (2001) provide a review. These laws refer to separate mortality characteristics for different age groups such as infants, youths and adults. When accounting for the dynamic development of mortality over time it seems natural to consider a factor model that accounts for these mortality laws. In this paper we assume the presence of multiple factors and impose structure on the loadings via specific functional forms.

The approach is similar to McNown and Rogers (1989). However, their model is both heavily over parametrised in terms of latent time-varying parameters, and does not fully exploit the information contained in the factor dynamics of the model.

The idea is similar to e.g. the dynamic Nelson-Siegel model for the term structure of interest rates, see Diebold and Li (2006). Diebold and Li suggest a factor model with parametrised factor loadings which identify level, slope and curvature of the yield curve, associated with the long, short and medium term yields. In the context of mortality we define loading functions that identify the factors that drive respectively infant mortality, adult mortality, and 'accident hump' (youth) mortality, plus a common factor uniformly affecting all age groups. We will generally refer to this class of factor models as parametric factor models (PFM). It follows from this approach that opposed to traditional factor analysis, the factors to be extracted will not necessarily be independent. In fact, the factors may potentially cointegrate when these are found to have unit roots.

We consider estimation of the model parameters and the factors by cross-section regressions over age groups for each period of time. These estimations are conducted over a grid of tuning parameters that define the shape of the loading functions and next a least squares criterion is used to determine the desired tuning parameters and the corresponding factor elements. This approach is similar to the first step of the cross section projection procedure suggested in Diebold and Li (2006). After the factors have been extracted the second step implies the estimation of a time series model for the factors. This can be done in different ways. For instance univariate as well as multivariate models for the factors can be formulated, and with the possibility of stationary as well as non-stationary factors that potentially cointegrate. The final time series specification of the factor dynamics is an empirical question and separate time series models are considered for this purpose.

We also consider a fully parametrised model specification formulated as a state space model. By use of the Kalman Filter recursions the model parameters and the factors can be estimated by full maximum likelihood.

We estimate the proposed model for women and men using French and US data for the sample period 1950-2014. The estimated functional forms appear to be rather similar across countries with the duration of the accident hump being longer for men than for women. The shape of the four factors also generally appear similar across countries but with differences across genders. In terms of model fit compared with the LC model, it appears that our model is doing especially well for explaining the age-specific death rates around the age groups defining the 'accident hump'.

We also evaluate the out of sample performance of the model where the predicted mortality rates are summarized in a loss function defined by the life expectancy. Specifically we apply the Model Confidence Set by Hansen et al. (2011) to evaluate the forecast performance on the horizons of 1, 10 and 20 years ahead with a number of benchmark models from the literature. We find that our model tends to be in the set of best predicting models for long horizon forecasts in particular.

We proceed as follows: In Section 2 we briefly describe the LC model and before providing a detailed description of the mortality data and set up a number of stylized facts of the mortality curve that a good mortality model should be able to reproduce in Section 3, similar to Diebold and Li (2006) for the yield curve. Section 4 introduce the PFM and its interpretation. Section 5 discusses the estimation of the PFM and in Section 6 we conduct the empirical analysis, describe the estimation results, the model fit, and the factor dynamics. Section 7 examines the out-of-sample forecast performance using the Model Confidence Set. Finally, Section 8 concludes.

2 The Lee-Carter Model

The observed data underlying the analysis are the death rates $m_{x,t}$ for age groups $x = 0, 1, \cdots, N$ at year $t = t_0, ...T$ and is defined as:

$$m_{x,t} = \frac{D_{x,t}}{ETR_{x,t}} \tag{1}$$

where $D_{x,t}$ is the number of deaths at age x in year t, and $ETR_{x,t}$, Exposure-to-Risk, is the population aged x in year t. In particular, $ETR_{x,t}$ can be interpreted as the average population size $P_{t,x}$ during year t at age x.

The Lee and Carter (1992) (LC) model describes the (log) age-specific death rates by:

$$\ln m_{x,t} = \alpha_x + \beta_x \kappa_t + \varepsilon_{x,t} \tag{2}$$

where α_x is a constant describing the general death rate for age group x. κ_t is a common time varying factor capturing the overall trend in death rates over time t. β_x is the factor loading capturing the effect of the factor κ_t on each age group x i.e. ages with high β_x will change more with κ_t. $\varepsilon_{x,t}$ is the age and year specific error not captured by the model. The LC model is basically a one-factor model that allows a common time trend to have age-specific loadings with respect to the development of the age-specific log death rates. Lee and Carter (1992) obtain identification using the normalizations $\sum_{x=0}^{N} \beta_x = 1$ and $\sum_{t=t_0}^{T} \kappa_t = 0$. The constraints imply that α_x becomes the time-average of the log death rates, $\ln m_{x,t}$, for each age x.[1] For estimating β_x and κ_t the singular value decomposition is applied to the matrix $(A)_{xt} = \ln m_{xt} - \alpha_x$ for all x, t. Lee and Carter (1992) find that κ_t can be modelled as a random walk with drift, although they allow for other specifications as well. The LC model is designed to maximize the in-sample fit by fitting a general factor model structure, using flexible factor loading, to the death rates. Note that the LC model does not impose any particular structure on the age-specific graduation of mortality; this is essentially data driven. However, it imposes a rigid structure on the improvements of the age-specific death rates over time by requiring these to be proportional and governed by the single factor κ_t.

3 Stylised Facts of the Mortality Curve

A good mortality model should desirably account for both the age (cross section) dimension of mortality as well as its development over time, i.e. the time dimension. Here, we describe some stylised facts of the (log) death rates to be modelled. For illustration, we use data for France and USA available through the Human Mortality Database (2015).[2] Figures 1a to 1d display log mortality on 10 year intervals from 1950 to 2010 for men and women.

[1] Following Nielsen and Nielsen (2014) the choice of restrictions is of no importance for the result of the likelihood function and forecasts. Other normalizations could be considered, however this gives an intuitive interpretation of α_x.

[2] In the empirical analysis For more information on the data source see Glei et al. (2014) and Andreeva and Barbieri (2016). We restrict the data to 1950 and onwards as this removes outliers, and we avoid structural changes in the exposure, see Lee and Miller (2001) and Booth et al. (2002). To avoid uncertainty about the death rates, due to few observations, we further restrict the ages to cover the ages 0 to 95 as is standard in the mortality forecasting literature.

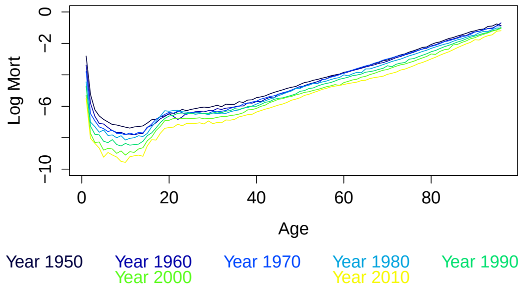

(a) Men, France

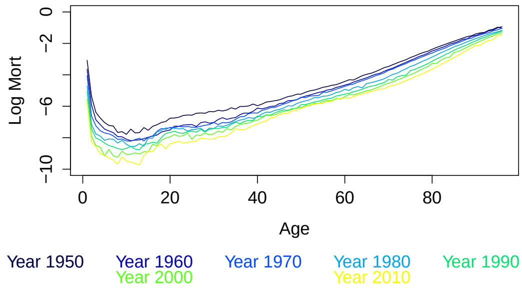

(b) Women, France

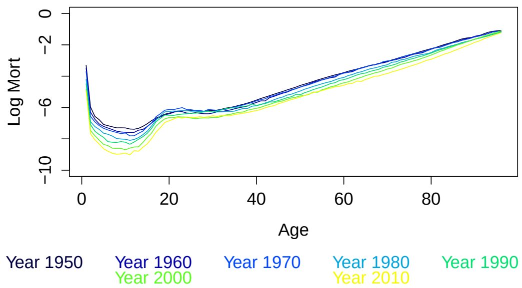

(c) Men, USA

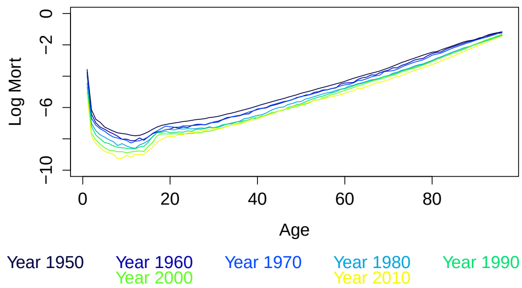

(d) Women, USA

Figure 1: The log age-specific death rates for the years 1950, 1960, 1970, 1980, 1990, 2000 and 2010 for men and women in France and the USA

The mortality curve shows a similar shape over the ages but the level of mortality tends to decline over time; the shape is very similar across both gender and countries. The infant mortality is seen to decline rapidly during early childhood. In the late teens the mortality rate experiences a rapid increase which is often termed the 'accident hump' which appears either as a distinct hump or as a flattening out of the death rates, see Heligman and Pollard (1980). After the accident hump, the mortality rates are gradually increasing almost linearly with age.

When investigating the time dynamics in the development of the log age-specific death rates, the review paper by Wong-Fupuy and Haberman (2004) notes that *"There is a broad consensus across the resulting projections: (1) an approximately log-linear relationship between mortality rates and time, (2) decreasing improvements according to age"*. The first point helps to explain the success of the LC model where the common time-varying factor is found to evolve almost linearly in most applications, see e.g. Lee and Miller (2001) and Callot et al. (2016). The log-linear time development of death rates are illustrated in Figures 2a to 2d. The second observation of decreasing improvements in mortality with respect to age can be described by the so called law of compensation mortality effect, see e.g. Gavrilov and Gavrilova (1991) or Thatcher (1999). In Figures 2a to 2d this effect is seen by a slope of the log mortality-time plot that decreases with age.

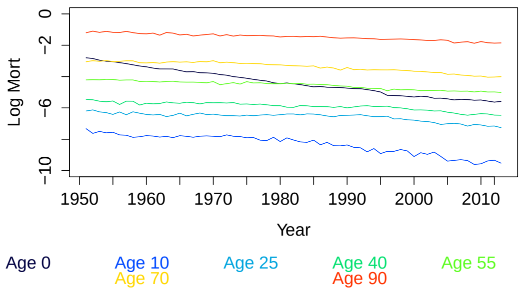

(a) Men, France

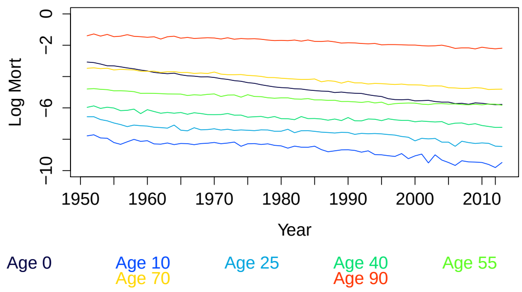

(b) Women, France

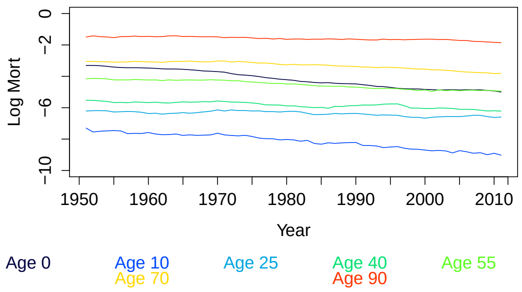

(c) Men, USA

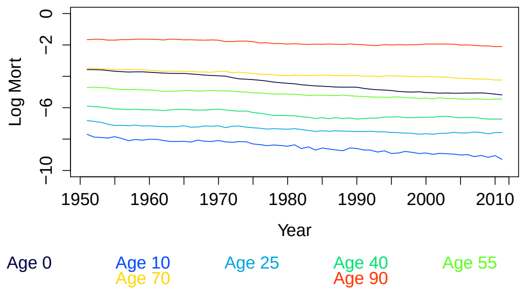

(d) Women, USA

Figure 2: The log age-specific death rates for a range of ages for French and American men and women from 1950-2012

Several studies find that a unit root is present in the individual age-specific death rates, see for instance Lazar and Denuit (2009). Also, it is common that the time-factor of the LC model is modelled as a random walk with drift. Basically this means that all death rates are governed by the same stochastic time trend component and hence for a system of $N+1$ age groups, all death rates cointegrate pairwisely and a total of N cointegrating relations exists among all age-specific death rates.

To examine this feature of the LC model we have conducted cointegration tests of all the pairwise combinations of log mortality across age groups. Note that due to the dimension of the data a full cointegration analysis cannot be conducted for the full data set. Figures 3a and 3d report a heatmap of the P-values from Johansens trace test, Johansen (1991), of the null hypothesis of zero cointegrating relations against one cointegrating relation for all combinations of the log age-specific death rates. As seen we cannot reject the null of no cointegration for most of the death rate pairs, especially for US data. It is apparent that most of the cointegrating relations found are between the adjacent ages, i.e. along the diagonal line. For both countries we find clear rejection of no cointegration amongst the youngest ages, but not for newborns. For France rejection of no cointegration among the oldest ages is found to a larger degree. Further it is found that around the accident hump and for the infants we cannot reject the null for relatively adjacent ages. Thus, overall the figures clearly show that the assumption of the LC model of N cointegrating relations is not consistent with the mortality data. The stochastic trends driving mortality over time are generally different across the age groups. The model we are subsequently going to propose will not have the restrictive feature of the LC model, since different factors are constructed to affect separate age groups.

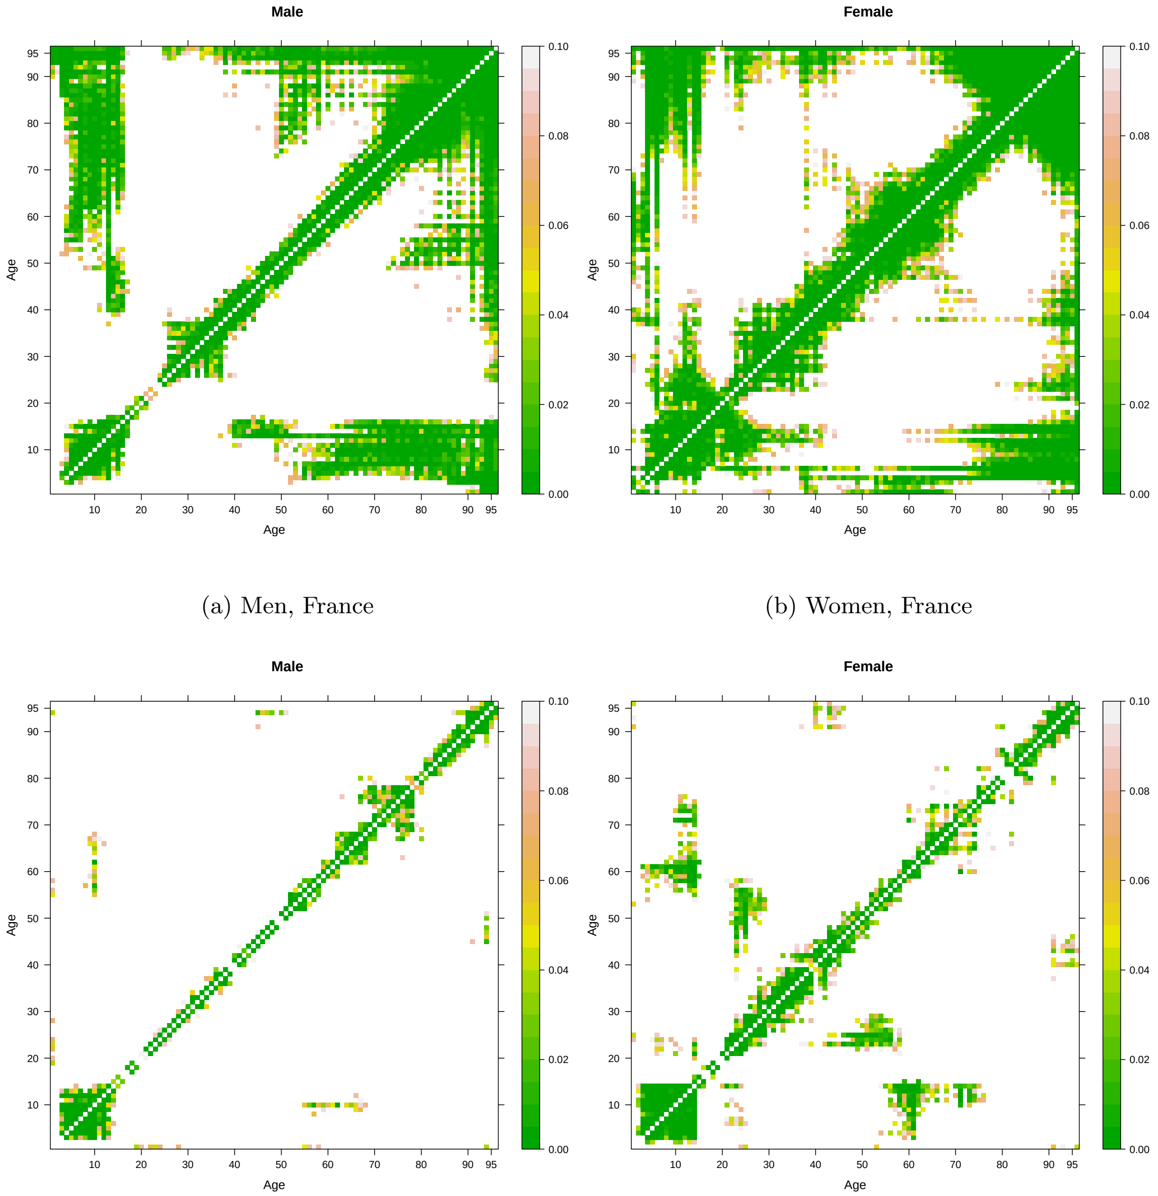

(a) Men, France (b) Women, France

(c) Men, USA (d) Women, USA

Figure 3: P-values from the Johansens Trace test for all pairwise combinations of the log age-specific death rates. The P-values are shown for significance levels between 0 and 0.10 for French and American men and women over the period 1950-2014. The test is performed with a restricted time trend in the cointegrating relation and 1 lag in the VAR specification. The P-values are obtained via the gamma aproximation from Doornik (1998); Doornik et al. (1999).

Summarizing, we observe six stylized facts for the term structure of (log) mortality: 1) declining mortality for infants, 2) increasing mortality around the accident hump, 3) linearly increasing mortality with age for adults, 4) a log-linear relationship between the death rates and time, 5) the log age-specific death rates are integrated of order one around a linear trend, 6) decreasing improvements in mortality with age, and 7) multiple stochastic trends characterize the development of log mortality over time for the different age groups.

4 The Parametric Factor Model for the Term Structure of Mortality

The model we are going to propose assumes that mortality is driven by multiple factors and we impose structure on the factor loadings capturing the regularities discussed in the previous section.

The proposed parametric factor (PFM) model reads as follows:

$$\ln m_{x,t} = \beta_{0,t} + \beta_{1,t}e^{-\lambda_1 x} + \beta_{2,t}e^{-\lambda_2(\log(x)-\log(c))^2} + \beta_{3t}\left(\frac{x}{N}\right)^{\lambda_3} + \varepsilon_{x,t} \tag{3}$$

The model has four factors $\beta_{it},\ i = 0,1,2,3$ with loading functions that are designed to capture distinct age groups. β_{0t} is a factor that is common to all age groups. The factor β_{1t} captures child mortality, β_{2t}, the accident hump, and finally, β_{3t} is a factor that tends to increase mortality with age. Note that the common factor has the constant loading one for all age groups. The loading for infant mortality declines rapidly with age. The loading for the accident hump is approximately bell-shaped around $\log(c)$ and finally the loading for the adult factor grows almost linearly with age when λ_3 is close to one. The loading functions estimated for France and the US are shown in Figures 4a and 4b; the estimation procedure will be discussed in the next section. Even though it may be claimed that the functional forms of the loading functions are arbitrary, they are designed such that the mortality laws and stylised facts described in Section 3 are captured through the model specification.

The level and infant terms $\beta_{0,t}$ and $\beta_{1t}e^{-\lambda_1 x}$, respectively, are used in many models using the age-specific graduation of mortality, see e.g. Siler (1979) and Rogers and Little (1994). The accident hump loading $e^{-\lambda_2(\log(x)-\log(c))^2}$ is taken from Heligman and Pollard (1980). The adult factor can be seen as a generalisation of the Gompertz model, similar to the Box and Cox (1964) power formulation. That is, the loading function captures the Gompertz specification if $\lambda_3 = 1$ and the Weibull (1951) model for λ_3 tending to zero.

5 Estimation Procedure for the Parametric Factor Model

We will consider two estimation procedures for estimating the PFM, the two-step procedure of Diebold and Li (2006) and exact maximum likehood estimation using the Kalman Filter recursions of the model written on state-space form similar to Diebold et al. (2006). Alternatively, one could use Bayesian estimation by Markov-Chain Monte Carlo methods or maximum likelihood estimation following Brouhns et al. (2002) assuming a Poisson distribution for the death counts.

5.1 The two step estimation procedure

The two step procedure considers first, to estimate the model parameters and the factors of the model, and second, estimating a time series model of the extracted factors with

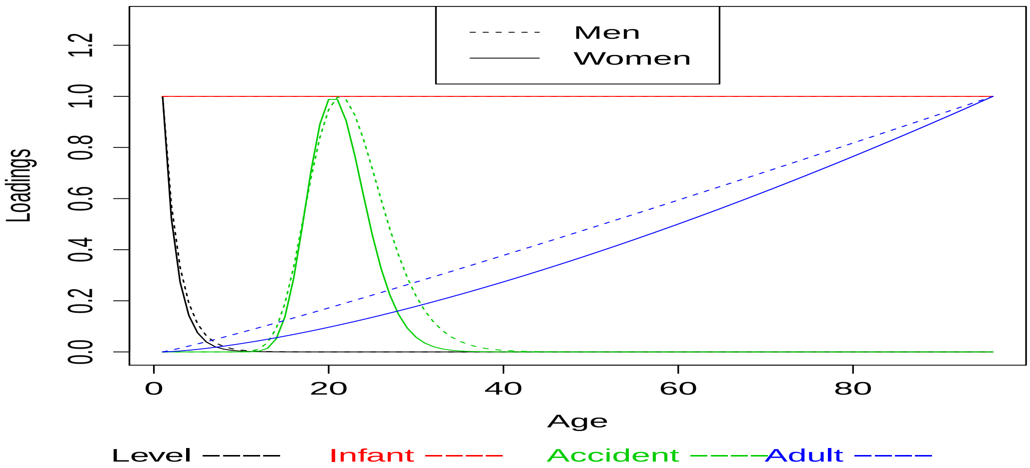

(a) Loadings, France

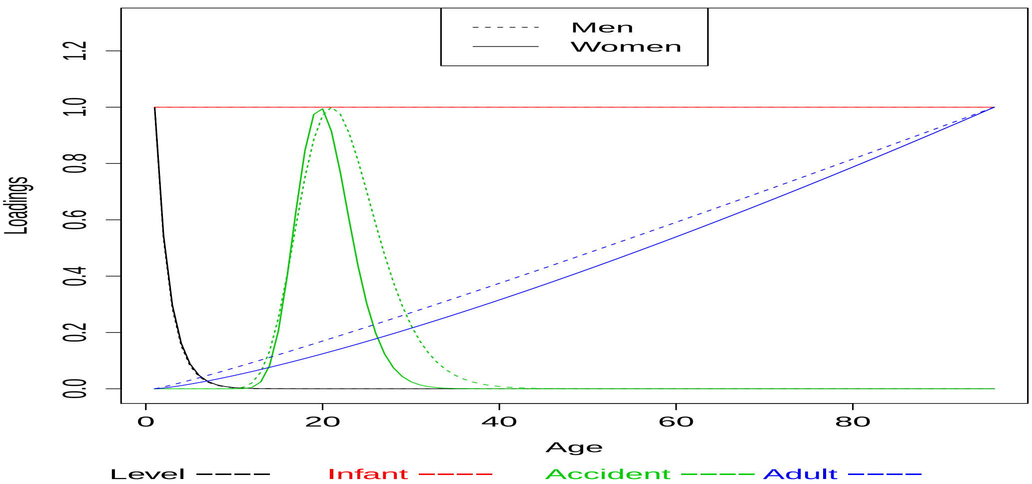

(b) Loadings, USA

Figure 4: Plot of the estimated loading functions for the years 1950-2012(2010) for men and women in France and USA. The loading functions correspond respectively to the level, infant, accident hump, and adult age groups. The loadings are estimated following the two-step procedure explained in section 5.

the primary purpose of forecasting. Regarding the first step, McNown and Rogers (1989) propose to estimate the factors by nonlinear least squares for each point in time. This allows not only the factors but also the model parameters to be time-varying. McNown and Rogers (1992) fix the parameters of the model a priori and estimate the factors in a sequence of cross-section regressions. The latter procedure is also the one adopted by Diebold and Li (2006) when estimating the dynamic Nelson-Siegel model for the term structure of interest rates, where the different loadings refer to the level, slope and curvature of the yield curve.

We suggest to modify McNown and Rogers (1992) and Diebold and Li (2006) by considering cross-sectional regressions for a fine grid of the model parameters and select the preferred model by the least squares criterion. This can also be implemented by a nonlinear least squares optimization algorithm such as the MaxBFGS procedure, which in fact is what we do in this paper. This step provides estimates of the four factors of the model. Note that as opposed to traditional factor models, the estimated factors will not be orthogonal and in fact are most likely to be dependent. In the second step of the two-step procedure time-series models are fitted to the factors. These can be univariate time series models such as ARIMA models, possibly with drifts or trends, or the factors can be modelled as stationary or nonstationary VAR models which potentially can allow for cointegration amongst the factors. It is an empirical question to properly select a time series model in the second step.

5.2 One-Step Estimation

The parametric factor model in Equation (3) can be formulated on state-space form and estimated by maximum likelihood by use of the Kalman Filter, see e.g. Durbin and Koopman (2012). This estimation procedure improves on the two-step esstimation procedure by allowing joint estimation of the latent factors and their transition dynamics as well as the unknown parameters $\lambda_1, \lambda_2, \lambda_3$ and c, assuming Gaussian errors. Estimating the system jointly delivers the correct inference compared with two-step approaches, which ignores the uncertainty and estimation errors from the first step in the second step.

The measurement equation of the state space model can be written as:

$$\ln \mathbf{m}_t = \Lambda \beta_t + \varepsilon_t$$

Where,

$$\ln \mathbf{m}_t = \begin{pmatrix} \ln m_{0,t} \\ \ln m_{1,t} \\ \vdots \\ \ln m_{N,t} \end{pmatrix}, \Lambda = \begin{pmatrix} 1 & e^{-\lambda_1 \cdot 0} & e^{-\lambda_2 (\log(0) - \log(c))^2} & \left(\frac{0}{N}\right)^{\lambda_3} \\ 1 & e^{-\lambda_1 \cdot 1} & e^{-\lambda_2 (\log(1) - \log(c))^2} & \left(\frac{1}{N}\right)^{\lambda_3} \\ \vdots & \vdots & \vdots & \vdots \\ 1 & e^{-\lambda_1 \cdot N} & e^{-\lambda_2 (\log(N) - \log(c))^2} & \left(\frac{N}{N}\right)^{\lambda_3} \end{pmatrix}, \beta_t = \begin{pmatrix} \beta_{0,t} \\ \beta_{1,t} \\ \beta_{2,t} \\ \beta_{3,t} \end{pmatrix} \quad (4)$$

$$\varepsilon_t = \begin{pmatrix} \varepsilon_{0,t} \\ \varepsilon_{1,t} \\ \vdots \\ \varepsilon_{N,t} \end{pmatrix} \quad (5)$$

where ε_t is assumed normal distributed as $N(0, I\sigma^2)$.

The transition equation of the state space model should be formulated to capture the dynamics of the factors. For instance if we assume that the factors are governed by a VAR(1) process in first differences, the transition equation can be specified as:

$$\begin{pmatrix} \beta_t \\ \Delta\beta_{t+1} \\ c \end{pmatrix} = \begin{bmatrix} I & I & 0 \\ 0 & \Phi & I \\ 0 & 0 & I \end{bmatrix} \begin{pmatrix} \beta_{t-1} \\ \Delta\beta_t \\ c \end{pmatrix} + \begin{pmatrix} 0 \\ v_t \\ 0 \end{pmatrix}$$

where v_t is multivariate normal distributed as $N(0, \Sigma)$.

In the case where the factors cointegrate with r cointegrating relations the transition dynamics can be written as:

$$\begin{pmatrix} \beta_t \\ \Delta\beta_{t+1} \\ c \end{pmatrix} = \begin{bmatrix} I & I & 0 \\ \alpha\beta' & \alpha\beta' & I \\ 0 & 0 & I \end{bmatrix} \begin{pmatrix} \beta_{t-1} \\ \Delta\beta_t \\ c \end{pmatrix} + \begin{pmatrix} 0 \\ v_t \\ 0 \end{pmatrix}$$

where v_t is multivariate normal distributed as $N(0, \Sigma)$. Where α is $4xr$ and β is $4xr$. Once again the first and third row are just mappings. The second row gives the desired VECM specification for the transition dynamics:

$$\Delta\beta_{t+1} = \alpha\beta' \left(\beta_{t-1} + \Delta\beta_t\right) + c + v_t \tag{6}$$

$$\Delta\beta_{t+1} = \alpha\beta'\beta_t + c + v_t \tag{7}$$

To estimate the parameters of the measurement and transition equations we use maximum likelihood on the prediction error decomposition, with the BFGS optimization algorithm. Estimation of the parameters $\psi = [\lambda_1, \lambda_2, \lambda_3, k, \sigma, \Sigma, \Phi$ (or α, β),] is then achieved via numerical optimization of the likelihood function given by the prediction error decomposition as

$$\mathcal{L}(\psi) = -\frac{NT}{2}\log 2\pi - \frac{1}{2}\sum_{t=1}^{T}\log|F_t| - \frac{1}{2}\sum_{t=1}^{T} v_t' F_t^{-1} v_t. \tag{8}$$

Thus, we note that the constant c is treated as a state parameter within the Kalman Filter. The calculations are implemented using the Ssfpack sofware package developed by Koopman et al. (1999) in the programming language Ox.

6 Empirical Analysis

6.1 Estimates using the two-step procedure

Figures 4a and 4b in Section 4 display the estimated shape of the loading functions for French and US men and women based on the two-step procedure. Table 1 reports the estimated shape parameters and their standard errors. The estimated parameters are similar across countries. However, the loading functions for the adult curvature is more convex for women than for men. Similarly, the shape and location of the accident hump varies across genders with men suffering from the accident hump longer than for women.

Table 1: Estimated loading function parameters and standard errors from the first step in the 2 step procedure for French and US men and women. The standard errors are calculated using the inverse Fisher information criterion

		Men					Women				
		λ_1	λ_2	λ_3	c	σ^2	λ_1	λ_2	λ_3	c	σ^2
Fr	Estimate	0.553	11.981	1.093	20.308	0.020	0.649	18.092	1.453	19.492	0.023
	Std. Err	0.013	0.024	0.004	0.002	0.018	0.013	0.060	0.003	0.005	0.018
US	Estimate	0.624	10.813	1.103	20.016	0.017	0.607	19.029	1.295	18.675	0.013
	Std. Err	0.013	0.020	0.003	0.002	0.018	0.010	0.042	0.003	0.004	0.018

The estimated factors are shown in Figures 5a to 5d for France and Figures 6a to 6d for the US.

A number of insights follow from these Figures. The factor governing the common mortality level decreases almost linearly and thus capturing a common decline in mortality across all age groups; this applies for both genders and countries. The infant factor for both men and women decline over the period showing that the infants have seen larger improvements in mortality reduction compared to the general level captured by the first factor. Moreover, it can be seen that the decline for the infant factor stagnates around 1995 for all populations considered. Hence, after 1995 the development in mortality for infants has generally followed the common rate.

The accident hump factor shows an increase in size from 1950 to about 1990 followed by stagnation for all but US men. Regarding the development of the adult factor Figures 5d and 6d exhibit an upward slope over the sample period and hence reducing the mortality improvements for the relevant age group. Thus, slower improvements in mortality with age is captured by the model, in line with the stylised facts.

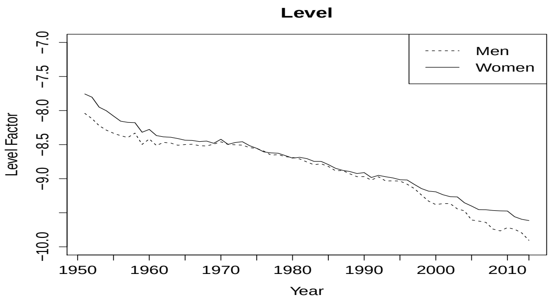

(a) The level factors

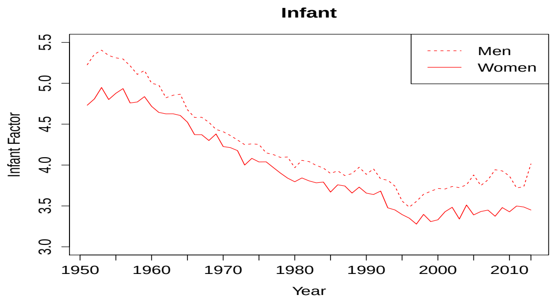

(b) The infant factors

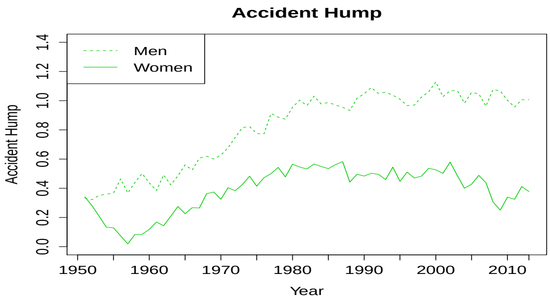

(c) The accident hump factors

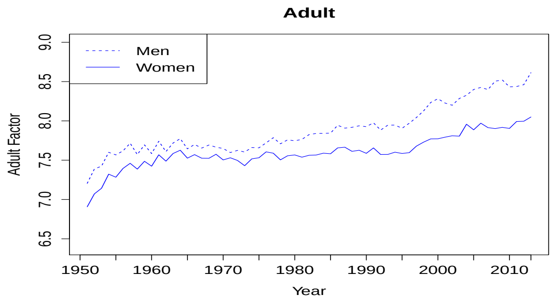

(d) The adult mortality factors

Figure 5: The factors β_{it}, $i = 0, 1, 2, 3$ are estimated by the two-step procedure for France using data from 1950-2014. The plots are showing respectively the level factor, infant factor, accident hump factor and adult factor for both genders.

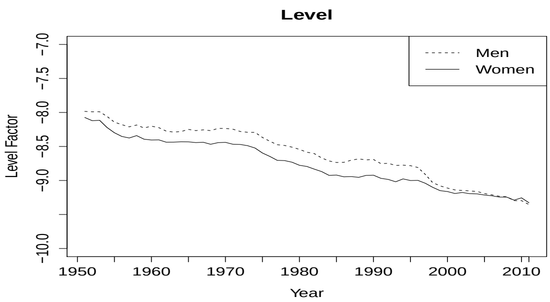

(a) The level factors

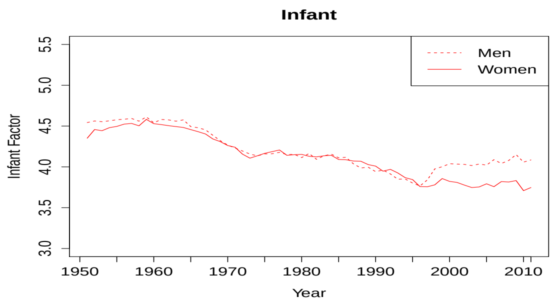

(b) The infant factors

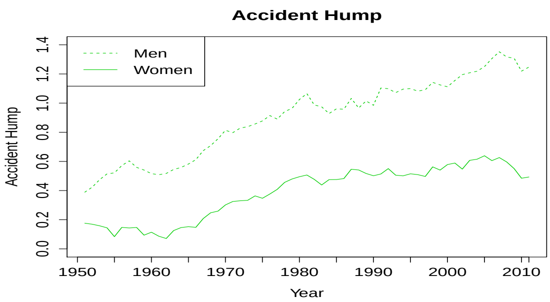

(c) The accident hump factors

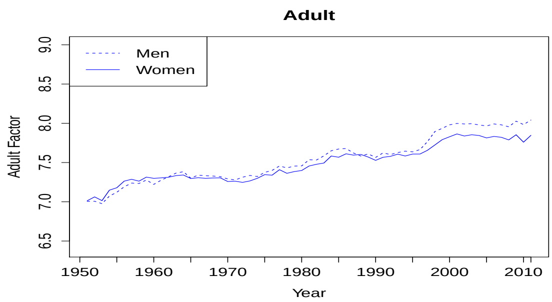

(d) The adult mortality factors

Figure 6: The factors β_{it}, $i = 0, 1, 2, 3$ are estimated by the two-step procedure for USA using data from 1950-2014. The plots are showing respectively the level factor, infant factor, accident hump factor and adult factor for both genders.

6.2 Cointegrating Analysis of the factors

In order to use the estimated model for forecast projections we need to examine the time series features of the estimated factors $\beta_{it}, i = 0, 1, 2, 3$. By using a range of unit root tests we find strong empirical support for the presence of unit roots, possibly with a drift, in all of the factors considered across both countries and gender. Given this observation it is not surprising that the age-specific log death rates individually appear to have similar time series characteristics. The one-factor model of Lee and Carter (1992) also typically model the factor as a random walk with drift. From visual inspection of the factors in Figures 5 and 6 it is evident that the various factors tend to co-move across gender and thus the factors are likely to cointegrated. Accounting for cointegration amongst the factors will potentially lead to superior forecasts.

We have conducted cointegration analysis using the Johansen trace test for different subsets of factors. In Table 2 we report the test results for each country and for each gender using all four factors. In Table 3 we examine tests for each country using all eight factors for both men and women, and finally, Table 4 displays the tests for men and women, respectively, by merging the factors across countries.

The results are rather different for the USA and France as can be seen from Table 2. For both genders the US factors are found not to cointegrate and hence these factors are driven by four separate common stochastic trends. On the other hand, the factors for French men and women appear to cointegrate with two or three cointegrating vectors and thus the factors for each gender appear to be driven by a single or possibly two common stochastic trends. This finding is also in line with the heat maps reported in Figure 3 showing that for France the pairwise log mortality rates appear more cointegrated compared to the USA.

In Table 3 the set of variables in the VAR model is expanded to include both men and women for each country. In this case the eight factors for the US data are driven by four common stocahstic trends. It is tempting to believe that the factors cointegrate across genders, however, a formal statistical test rejects this hypothesis. For the French data the eight factors have six to seven cointegrating vectors and thus have one or two common stochastic trends. Again, a formal test rejects that the factors cointegrate pairwisely across genders.

Finally, Table 4 shows that when pooling the US and French data for men and women respectively, both the male and female factors are likely to be driven by 6 factors and thus have two cointegrating relations. Hence cross-country similarities exist across countries for both genders but only to a limited extent.

These findings demonstrate that different time series specifications should be considered when modelling the factors with the purpose of forecasting. For the US it seems appropriate to specify a VAR in first differences with a vector of unresticted constants to capture the drift of the single series. It could also be considered to base predictions on an expanded (cointegrated) VAR model including factors for both genders. For France, a cointegrated VAR with cointegration rank two or three seems appropriate. An expanded cointegrated VAR with eight factors and six to seven cointegrating vectors is also possible. When modelling the factors as univariate time series models a random walk with drift specification is appropriate but since the cross dependence of factors is neglected in this case, it is likely that inferior forecasts will result.

Table 2: Test for cointegration rank amongst factors for US and French men and women.

	USA				France			
	Men		Women		Men		Women	
Rank	Trace-test	p-value	Trace-test	p-value	Trace-test	p-value	Trace-test	p-value
0	52.240	[0.323]	47.939	[0.512]	106.790	[0.000]**	108.730	[0.000]**
1	27.748	[0.641]	30.689	[0.468]	56.044	[0.001]**	47.599	[0.014]*
2	13.926	[0.668]	15.243	[0.562]	28.089	[0.024]*	24.893	[0.064]
3	1.1522	[0.992]	3.0677	[0.858]	3.642	[[0.788]	9.2236	[0.171]

*Note: The Johansen trace test is calculated with a trend restricted to the cointegration space. The number of lags in the VAR is 1 for all cases based on HQ, SIC and Portmanteau tests. "**" and "*" signify significance at the 1 % and 5% level, respectively.*

Table 3: Test for cointegration rank amongst factors for men and women for USA and France.

	USA		France	
	Men and Women		Men and Women	
Rank	Trace-test	p-value	Trace-test	p-value
0	286.700	[0.000]**	287.730	[0.000]**
1	194.530	[0.000]**	215.400	[0.000]**
2	139.460	[0.001]**	158.870	[0.000]**
3	89.751	[0.041]*	116.050	[[0.000]**
4	52.267	[0.321]	76.354	[0.002]**
5	34.795	[0.257]	47.410	[0.015]*
6	19.806	[0.240]	24.016	[0.082]
7	7.897	[0.268]	8.524	[0.218]

*Note: The Johansen trace test is calculated with a trend restricted to the cointegration space. The number of lags in the VAR is 1 for all cases based on HQ, SIC and Portmanteau tests. "**" and "*" signify signifcance at the 1 % and 5% level, respectively.*

Table 4: Test for cointegration rank amongst factors for US and French men and women.

	Men		Women	
	USA and France		USA and France	
Rank	Trace-test	p-value	Trace-test	p-value
0	225.130	[0.000]**	221.980	[0.000]**
1	158.190	[0.016]*	152.790	[0.036]*
2	111.620	[0.113]	107.810	[0.178]
3	75.429	[0.311]	70.517	[[0.490]
4	47.924	[0.513]	44.325	[0.679]
5	27.275	[0.668]	27.285	[0.667]
6	11.364	[0.850]	13.676	[0.688]
7	3.867	[0.759]	3.686	[0.783]

*Note: The Johansen trace test is calculated with a trend restricted to the cointegration space. The number of lags in the VAR is 1 for all cases based on HQ, SIC and Portmanteau tests. "**" and "*" signify signifcance at the 1 % and 5% level, respectively.*

6.3 Estimates using the one-step procedure

We now consider one step estimation of the model employing maximum likelihood estimation via the Kalman Filter recursions with the model specified on state space form. This method theoretically improves the efficiency as it avoids the issue in the two-step estimator of ignoring the estimation error from the first step in the second step. The estimation for

US is based on the assumption of VAR(1) in first difference for the transition dynamics and for France it is based on the cointegrated VAR model with 2 cointegrating relations (The results for France are to appear later). These dynamics where chosen as they where found to appropriately describe the transition dynamics in Section 6.2. Table 5 reports the estimated shape parameters and their standard errors for using the state space approach. It is seen that the loading parameters and their standard errors are very similar to those obtained from the two-step procedure. However, it is found that the standard error for both the shape and location of the accident hump are larger than for the two-step approach. In contrast to what would be expected. As the parameters are almost identical to the two-step approach we omit plotting the factor loadings which will be similar to those displayed in Figure 4.

Table 5: Estimated loading function parameters and standard errors from the one-step procedure for French and US men and women. For US the VAR(1) model in first difference is assumed for the transition dynamics and for France a VECM with 2 cointegrating relations is assumed. The standard errors are calculated using the inverse Fisher information criterion

		Men					Women				
		λ_1	λ_2	λ_3	c	σ^2	λ_1	λ_2	λ_3	c	σ^2
Fr	Estimate	-	-	-	-	-	-	-	-	-	-
	Std. Err	-	-	-	-	-	-	-	-	-	-
US	Estimate	0.630	10.248	1.101	20.204	0.019	0.615	17.307	1.290	18.915	0.014
	Std. Err	0.013	0.113	0.003	0.038	0.018	0.011	0.382	0.003	0.064	0.018

Figure 7 show the estimated factors (or states) for the one-step state space estimation procedure for US based on the VAR(1) specification in differences (The factors estimates for the VECM for France are still to appear).

Comparing the estimated factors with those obtained in the first step of the two-step approach the results appear similar. Note however, that the factors from the one-step estimation show a smoother development over the years. This is because the one-step procedure directly accounts for the transition dynamics in the joint estimation.

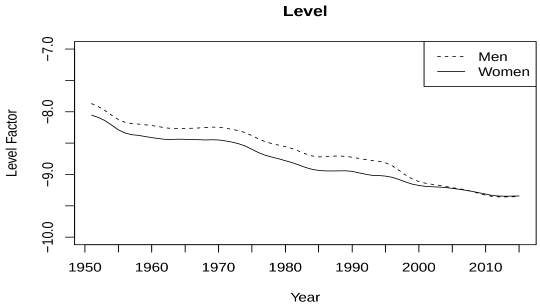

(a) The level factors

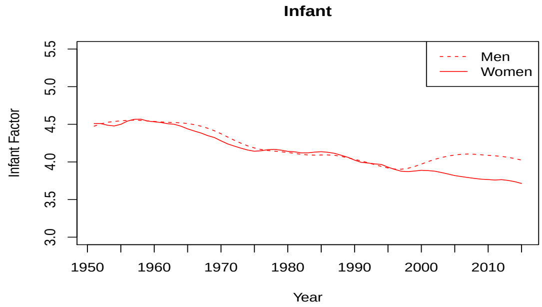

(b) The infant factors

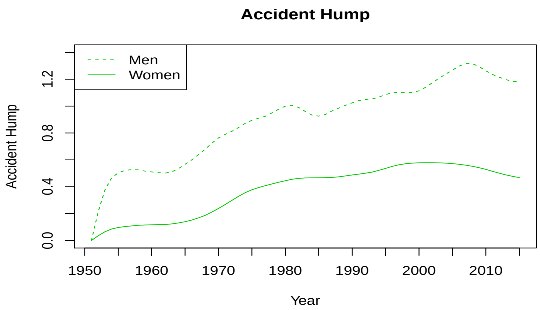

(c) The accident hump factors

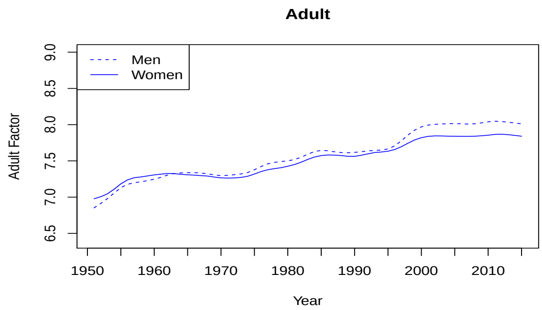

(d) The adult mortality factors

Figure 7: The factors $\beta_{it},\ i = 0, 1, 2, 3$ are estimated using the one-step procedure for USA using data from 1950-2014 and assuming a VAR(1) dynamics for the first difference of the states. The plots are showing respectively the level factor, infant factor, accident hump factor and adult factor for both genders.

6.4 Model Fit

We now compare the PFM with the Lee-Carter model in terms of in-sample fit.

As the PFM does not include a constant for each age-specific death rate we are interested in whether the model can capture the mean by relatively few parameters. As seen in Figures 8a to 8d, the model captures the mean well for all populations. Note that by construction α_x in the LC model is equal to the mean of the age specific log death rates, which corresponds to the data mean levels in the Figures.

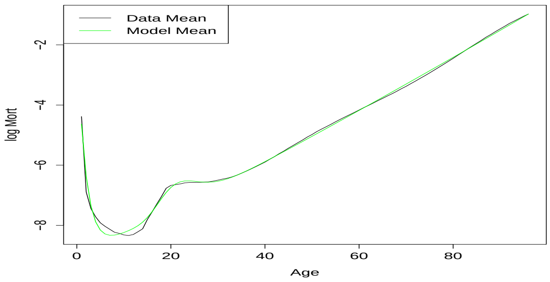

(a) Men, France

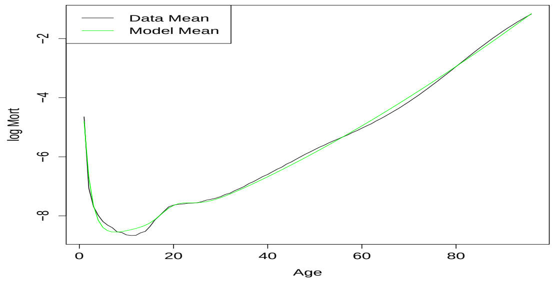

(b) Women, France

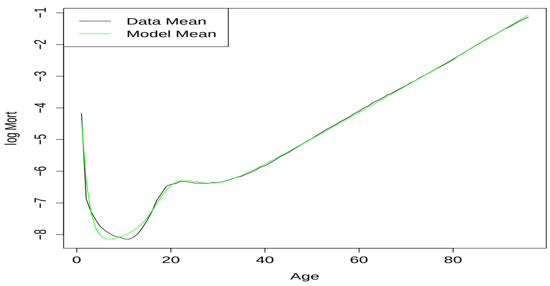

(c) Men, USA

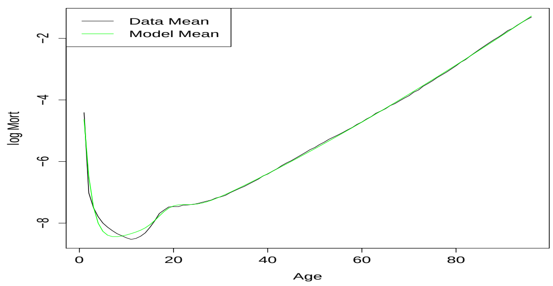

(d) Women, USA

Figure 8: The mean of the data and the mean of the parametric factor model for both men and women, for France and USA. The estimation period is 1950-2012.

To further quantify the model fit, we calculate a pseudo R^2 for each age group by running a regression of the age-specific death rates on a constant and the fitted values.[3] The pseudo R^2's shown in Figures 9a to 9d display that both the LC and the PFM fit the observed data well. However, the PFM tends to perform better around the accident hump, where the LC model is found to have poor performance.

[3]This corresponds to the partial correlation squared between the fitted and observed values.

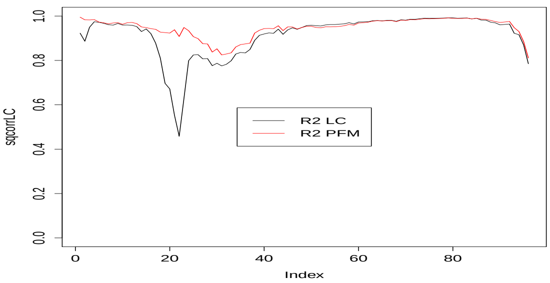

(a) Men, France

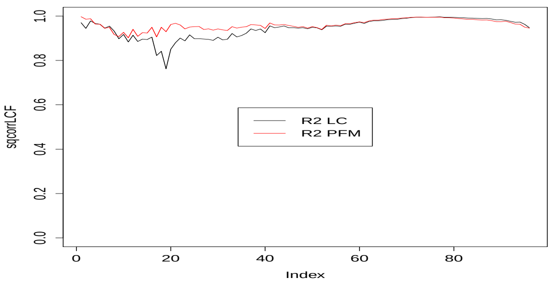

(b) Women, France

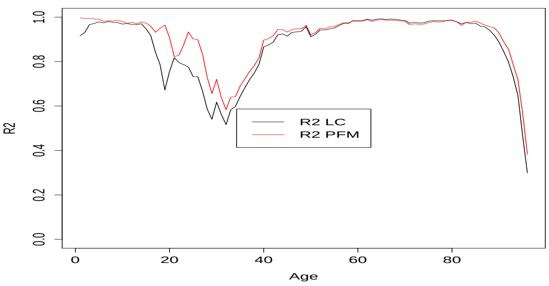

(c) Men, USA

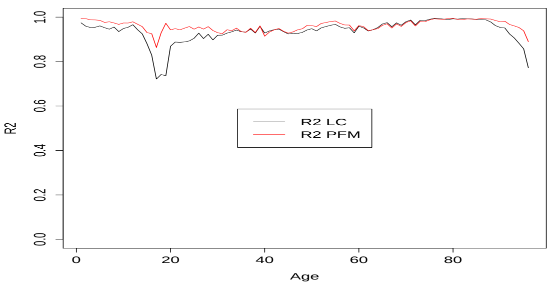

(d) Women, USA

Figure 9: The pseudo R^2 (within) for the PFM and the LC model for all ages. The R^2 is shown for both men and women in France and the USA respectively.

Next, we investigate how each of the factors contributes to the explanatory power of the model by calculating the partial correlation between the log mortality and a particular factor after adjusting for the influence of the fit obtained from the remaining factors. This adjustment is necessary because the factors are non-orthogonal. The Figures 10a to 10d display the partial correlations in excess of a 65% threshold for all ages to identify where the different factors improve the fit.

It is seen that the infant mortality factor significantly improves the fit for infants as desired. The level factor substantially improves the performance for most ages, and the accident hump factor primarily affects the mortality in the years around the accident hump. Finally, the adult factor primarily improves the fit for the adult ages as desired, but its partial explanatory power is of a smaller magnitude compared with the other factors, mainly because the adult factor is highly correlated with the factor common to all age groups.

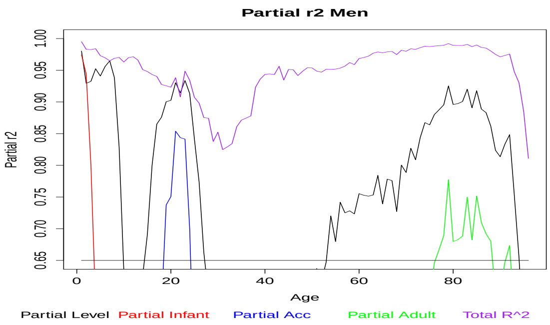

(a) Men, France

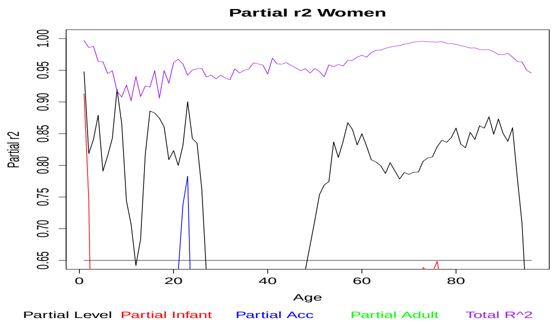

(b) Women, France

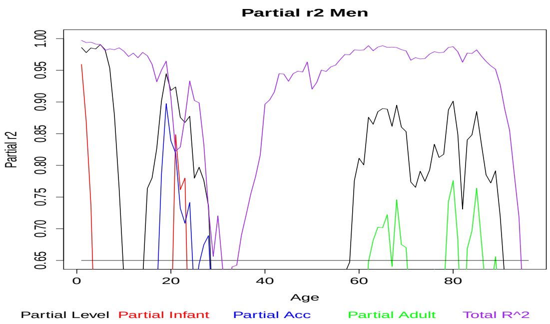

(c) Men, USA

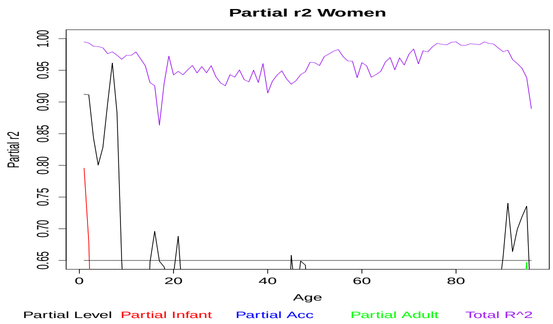

(d) Women, USA

Figure 10: The partial r^2 for the infant, level, accident hump, and adult factor respectively. The relative improvement from each of the factors are shown in excess of a 65% threshold. This is shown for both genders and France and the USA respectively.

7 Forecast Evaluation

In this section we investigate the forecast performance of the PFM and compare with relevant benchmark models. For forecast evaluation and comparison we use the Model Confidence Set (MCS) approach developed of Hansen et al. (2011).[4]

The MCS procedure is a test for predictive ability across a number of competing models, which sequentially removes the model that performs significantly worse than the remaining models left in the model confidence set. The procedure delineates the set of best performing models at a given confidence level among which we cannot say that any of the other models perform statistically better.

Hence, the MCS does not necessarily pick out a single best model but rather delineates a set of best models as the available information might not be able to discriminate between these models. The MCS procedure returns p-values, $\hat{p}_i$, for each model i considered, and from this the MCS can be determined. The MCS procedure returns a p-value of 1 to the best performing model.[5]

To reduce the dimension of the forecast evaluation we use the life expectancy which aggregates the forecasted age-specific death rates into a single measure. The life expectancy is

[4] The MCS approach is implemented via the Ox-package Mulcom 3.0 by Hansen and Lunde (2014)

[5] For the case with only two models the forecast performance could be tested via the Diebold and Mariano (1995) test, which only allows for pairwise comparisons, whereas the MCS procedure allows for joing multiple model evaluation.

calculated as in Brouhns et al. (2002) assuming a piecewise constant hazard rate given by:

$$\bar{e}_x^{\uparrow}(t) = \frac{1-\exp\left(-m_{x,t}\right)}{m_{x,t}} + \sum_{k\geq 1}^{max(x)} \left(\prod_{j=0}^{k-1} \exp\left(-m_{x+j,t}\right) \right) \frac{1-\exp\left(-m_{x+k,t}\right)}{m_{x+k,t}} \tag{9}$$

where $\bar{e}_x^{\uparrow}(t)$ is the period life expectancy, and $m_{x,t}$ signifies the age-specific death rates.

To show the robustness of the proposed model at producing reliable forecasts we consider data for men and women for the USA and France in the forecast evaluation. The forecasts are constructed by recursively estimating each model from 1950 onwards until year $t = 1970, 1971, \ldots$ and forecasting 1, 10, and 20 years ahead. This gives respectively 43, 34, and 24 forecasts of the age-specific death rates for each model. The forecast performance is evaluated using the mean squared error of the life expectancy as the loss function. For implementation we use the block bootstrap with a block length of 2 and a confidence level of 5%, see Hansen et al. (2011) for details.

As benchmark models we use 1) a random walk with drift specification for each (log) age specific death rate, 2) the Lee and Carter (1992) model with a single factor, 3) and the functional data approach (FDA) of Hyndman et al. (2007). Based on the analysis in Section 5 and 6 we consider two dynamic specifications of the factor structure, a VECM (with two cointegrating relations) and a VAR(1) in first differences of the factors. For comparisons we use both specifications for each country and gender estimated by the two-step procedure. For the one-step procedure we consider estimation assuming the VECM structure for France and the VAR(1) structure in differences for the US (Results still to appear!). Using the two-step procedure we further compare a VAR(1) model in levels and univariate ARIMA model.[6]

For the LC model we use a random walk with drift specification for the single factor κ_t. The FDA model of Hyndman et al. (2007) can be considered an extension of the LC model by using K factors and smoothing across the death rates. [7] The results are reported in Table 6 for France and in Table 7 for the USA.

Table 6: Forecasting life expectancy 1, 10, and 20 years ahead with mean-squared error criterion for French men and evaluated using the Model Confidence Set. Mean squared error along with p-values for the estimated model confidence set for life expectancy. The models in the set of best models are denoted with a ∗.

France	Men						Women					
	1 Year		10 Year		20 Year		1 Year		10 Year		20 Year	
Model	MSE	P-val	MSE	P-val	MSE	P-val	MSE	P-val	MSE	P-val	MSE	P-val
PFM VAR(1)	0.098	0.000	2.493	0.005	13.480	0.000	0.164	0.000	1.719	0.000	7.529	0.000
PFM ARIMA’s	0.103	0.001	0.968	0.856*	4.147	0.906*	0.071	0.002	0.144	0.055*	0.370	0.920*
PFM ΔVAR(1)	0.095	0.007	0.985	0.856*	4.611	0.182*	0.086	0.001	0.157	0.008	0.444	0.320*
PFM VECM	0.109	0.006	0.955	1.00*	4.121	1.00*	0.085	0.001	0.250	0.000	0.835	0.068*
PFM KF - VECM	-	-	-	-	-	-	-	-	-	-	-	-
RW w. drift	0.032	0.638*	1.135	0.436*	5.439	0.001	0.032	1.00*	0.103	1.00*	0.367	1.00*
LC	0.099	0.001	1.479	0.001	6.367	0.000	0.119	0.000	0.229	0.008	0.460	0.254*
FDA	0.030	1.00*	1.085	0.840*	5.436	0.022	0.037	0.289*	0.410	0.053*	1.581	0.254*

[6]These specifications have often been used in studies applying graduation laws of mortality, see Booth and Tickle (2008); McNown and Rogers (1989, 1992)

[7]The factors are estimated using weighted principal components in the R package Demography, see Hyndman et al. (2007) and Hyndman et al. (2011) for further details.

Table 7: Forecasting life expectancy 1, 10, and 20 years ahead with mean-squared error criterion for US men and women evaluated using the Model Confidence Set. Mean squared error along with p-values for the estimated model confidence set for life expectancy. The models in the set of best models are denoted with a $*$.

USA	Men						Women					
	1 Year		10 Year		20 Year		1 Year		10 Year		20 Year	
Model	MSE	P-val	MSE	P-val	MSE	P-val	MSE	P-val	MSE	P-val	MSE	P-val
PFM VAR(1)	0.113	0.007	1.607	0.014	5.915	0.000	0.069	0.004	1.515	0.049	11.630	0.038
PFM ARIMA's	0.112	0.008	0.787	1.00*	2.493	0.283*	0.054	0.008	0.519	0.124*	1.378	0.031
PFM ΔVAR(1)	0.110	0.005	0.897	0.275*	3.094	0.004	0.057	0.008	0.562	0.004	1.369	0.195*
PFM VECM	0.122	0.003	1.205	0.275*	2.159	1.00*	0.064	0.008	0.714	0.016	1.860	0.070*
PFM KF-ΔVAR(1)	-	-	-	-	-	-	-	-	-	-	-	-
RW w. drift	0.035	1.00*	1.240	0.135*	4.016	0.000	0.023	1.00*	0.435	1.00*	1.243	1.00*
LC	0.138	0.007	1.771	0.015	5.375	0.000	0.080	0.006	0.682	0.033	1.790	0.009
FDA	0.044	0.127*	1.577	0.122*	4.824	0.001	0.026	0.032	0.495	0.148*	1.441	0.070*

France, men. For French men the MCS using a 1-year forecast horizon includes the RW w. drift and FDA specifications. However, when expanding the forecast horizon the MCS now includes 3 variants of the PFM and in fact for a twenty year forecast horizon the MCS excludes the RW and FDA specifications. It is interesting to observed that in this forecast competition the LC model is never included in the MCS. The same applies for the PFM model specification where the factors are modelled as a VAR(1) in levels. This is not surprising because all the factors were found to have unit roots.

USA, men. The pattern observed for French men generally applies for US men as well. However, for the 20 year horizon only two of the PFM models are included in the MCS.

France, women. For French women and a forecasting horizon of 1 year the results are rather similar to those of French men and in particular the RW w. drift and the FDA model are the ones included in the MCS. For a 10 year horizon the MCS also includes a single PFM specification and for a 20 year horizon only the PFM with a VAR(1) in levels is not included in the MCS.

USA, women. For US women the RW with drift model is always in the MCS. For a 10 year horizon the results are similar to French women and for a 20 year horizon the MCS is slightly smaller than for French women and includes in particular the two PFM specifications the FDA and the RW w. drift specifications.

Summarizing, the PFM class of models appears to perform especially well for longer forecast horizons and in most cases performs better than the LC model. An explanation for this result could be the structural features of the PFM class of models compared to the LC model. For longer horizons the structural restrictions on the loadings account for different factors affecting the separate age groups. The structure implied by the PFM specification ensure a realistic shape of the mortality curve, which cannot be captured by a single factor LC model. Another conclusion is that in situations where competing models are performing well, especially for longer horizons, the different PFM models also perform well. On the other hand in situations where competing models are not performing so well, the PFM models are included in the MCS as seen especially for men.

8 Concluding Remarks

We have suggested a multi-factor model for the term structure of mortality. The factors are identified after restrictions on the loading functions in such a way that different age groups and their factor dynamics can be addressed separately. So rather than having a single factor governing all age groups as seen for the LC model, different factors (or trends) play a role in the way that mortality across age groups develop. In particular, we consider seperate factors driving infant mortality, the accident hump mortality, mortality for the elderly in addition to a common factor affecting all age groups. We have suggested 2 estimation methods that are similar to estimation of term structure models considered in other contexts. In an application we apply the methodology to mortality data for the US and France for each gender. The models are shown to provide good fit and for certain age groups a better fit compared to the LC model. In a forecast comparison across a range of competing models the new class of models that we consider in the paper are shown to perform well, especially over longer forecast horizons.

References

ANDREEVA, M. AND M. BARBIERI (2016): "ABOUT MORTALITY DATA FOR UNITED STATES," Tech. rep., http://www.mortality.org/hmd/FRATNP/InputDB/FRATNPcom.pdf.

BOOTH, H., J. MAINDONALD, AND L. SMITH. (2002): "Applying Lee-Carter under conditions of variable mortality decline," *Population Studies*, 56, 325–336.

BOOTH, H. AND L. TICKLE (2008): "Mortality modelling and forecasting: A review of methods," *Annals of Actuarial Science*, 3, 3–43.

BOX, G. E. AND D. R. COX (1964): "An analysis of transformations," *Journal of the Royal Statistical Society. Series B (Methodological)*, 211–252.

BROUHNS, N., M. DENUIT, AND J. K. VERMUNT (2002): "A Poisson log-bilinear regression approach to the construction of projected lifetables," *Insurance: Mathematics and Economics*, 31, 373–393.

CALLOT, L., N. HALDRUP, AND M. KALLESTRUP-LAMB (2016): "Deterministic and stochastic trends in the Lee–Carter mortality model," *Applied Economics Letters*, 23, 486–493.

CZADO, C., A. DELWARDE, AND M. DENUIT (2005): "Bayesian Poisson log-bilinear mortality projections," *Insurance: Mathematics and Economics*, 36, 260–284.

DATABASE, H. M. (2015): "Human Mortality Database." University of California, Berkeley (USA), and Max Planck Institute for Demographic Research (Germany), URL: www.mortality.org, downloaded 2015.

DE JONG, P. AND L. TICKLE (2006): "Extending Lee–Carter mortality forecasting," *Mathematical Population Studies*, 13, 1–18.

DIEBOLD, F. AND R. MARIANO (1995): "Comparing Predictive Accuracy," *Journal of Business and Economic Statistics*, 13, 253–263.

DIEBOLD, F. X. AND C. LI (2006): "Forecasting the term structure of government bond yields," *Journal of econometrics*, 130, 337–364.

DIEBOLD, F. X., G. D. RUDEBUSCH, AND S. B. ARUOBA (2006): "The macroeconomy and the yield curve: a dynamic latent factor approach," *Journal of econometrics*, 131, 309–338.

DOORNIK, J. A. (1998): "Approximations to the asymptotic distributions of cointegration tests," *Journal of Economic Surveys*, 12, 573–593.

DOORNIK, J. A. ET AL. (1999): "Erratum [Approximations to the Asymptotic Distribution of Cointegration Tests]," *Journal of Economic Surveys*, 13.

DURBIN, J. AND S. J. KOOPMAN (2012): *Time series analysis by state space methods*, 38, Oxford University Press.

GAVRILOV, L. A. AND N. S. GAVRILOVA (1991): "The biology of life span: a quantitative approach." .

Girosi, F. and G. King (2008): *Demographic forecasting*, Princeton University Press.

Glei, D., F. Mesl, J. Vallin, J. Wilmoth, and M. Barbieri (2014): "ABOUT MORTALITY DATA FOR FRANCE, TOTAL POPULATION," Tech. rep., http://www.mortality.org/hmd/FRATNP/InputDB/FRATNPcom.pdf.

Gompertz, B. (1825): "On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies," *Philosophical transactions of the Royal Society of London*, 513–583.

Hansen, P. and A. Lunde (2014): "MulCom 3.00, an Ox software package for multiple comparisons," .

Hansen, P. R., A. Lunde, and J. M. Nason (2011): "The model confidence set," *Econometrica*, 79, 453–497.

Heligman, L. and J. H. Pollard (1980): "The age pattern of mortality," *Journal of the Institute of Actuaries*, 107, 49–80.

Hyndman, R. J., H. Booth, L. Tickle, and J. Maindonald (2011): *Hyndman R J, ed. (2011), demography: Forecasting mortality and fertility data , URL: http://www.robhyndman.info/Rlibrary/demography, R package.*

Hyndman, R. J., S. Ullah, et al. (2007): "Robust forecasting of mortality and fertility rates: a functional data approach," *Computational Statistics & Data Analysis*, 51, 4942–4956.

Johansen, S. (1991): "Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models," *Econometrica: Journal of the Econometric Society*, 1551–1580.

Koopman, S. J., N. Shephard, and J. A. Doornik (1999): "Statistical algorithms for models in state space using SsfPack 2.2," *The Econometrics Journal*, 107–160.

Lazar, D. and M. M. Denuit (2009): "A multivariate time series approach to projected life tables," *Applied Stochastic Models in Business and Industry*, 25, 806–823.

Lee, R. D. and L. R. Carter (1992): "Modeling and Forecasting U.S. Mortality," *Journal of the American Statistical Association*, 87, 659–671.

Lee, R. D. and T. Miller (2001): "Evaluating the performance of the Lee-Carter method for forecasting mortality." *Demooraphy*, 38, 537–549.

Makeham, W. M. (1860): "On the law of mortality and the construction of annuity tables," *The Assurance Magazine, and Journal of the Institute of Actuaries*, 301–310.

McNown, R. and A. Rogers (1989): "Forecasting mortality: A parameterized time series approach," *Demography*, 26, 645–660.

——— (1992): "Forecasting cause-specific mortality using time series methods," *International Journal of Forecasting*, 8, 413–432.

Nielsen, B. and J. P. Nielsen (2014): "Identification and forecasting in mortality models," *The Scientific World Journal*, 2014.

Renshaw, A. E. and S. Haberman (2003): "Lee–Carter mortality forecasting with age-specific enhancement," *Insurance: Mathematics and Economics*, 33, 255–272.

Rogers, A. and J. S. Little (1994): "Parameterizing age patterns of demographic rates with the multiexponential model schedule," *Mathematical Population Studies*, 4, 175–195.

Siler, W. (1979): "A competing-risk model for animal mortality," *Ecology*, 750–757.

Tabeau, E., A. van den Berg Jeths, and C. Heathcote (2001): *Forecasting mortality in developed countries: Insights from a statistical, demographic and epidemiological perspective*, vol. 9, Springer Science & Business Media.

Thatcher, A. R. (1999): "The long-term pattern of adult mortality and the highest attained age," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 162, 5–43.

Weibull, W. (1951): "Wide applicability," *Journal of applied mechanics.*

Wong-Fupuy, C. and S. Haberman (2004): "Projecting mortality trends: recent developments in the United Kingdom and the United States," *North American Actuarial Journal*, 8, 56–83.